

Sept. 14, 2021



Internet History Webinar

How the US NSFNET Regional Academic Networks Evolved
to Become the Global Internet (1985 — 1995)



Internet2 Internet History Webinar: How the US NSFNET Regional Academic Networks Evolved to Become the Global Internet (1985 — 1995) – 14 September 2021

Ana Hunsinger

Thank you. Welcome everyone. Hope you're doing well. Before we begin just a few housekeeping announcements: this event is being recorded, it will be available to you in the community after the session. We also plan to take questions through the q&a button at the bottom of your screen, we will try our best to monitor the chat window so you can also post questions or comments there as well. We'll get started in just a few seconds.

Good day everyone, and thank you for joining us from wherever you may be located. My name is Ana Hunsinger, and I'm the Vice President of Community Engagement at Internet2. I am so happy to be here today and to welcome you to today's Internet history webinar. Alongside the unique Dennis Jennings, Internet2 is pleased to help host and organize this webinar, where we have the opportunity to come together and celebrate the organizations and the individuals who helped create the global Internet.

This is our second history webinar, and what we hope will be an ongoing series where we will have the opportunity to share with each other, the many stories about the rich history of the development of the Internet, and the many people that played a critical role and that helped

shape the critical infrastructure and services especially for the research and education community. Last year, we had the opportunity to help celebrate the 35th anniversary of the National Science Foundation network, also known as the NSFNET. Many of you know that NSFNET was a pivotal program, and a collaboration that laid the foundation for our nation's Research and Education networks, and the global Internet as we know it today. As this year 2022 marks the 25th anniversary of Internet2 and this milestone would not have been possible if not for the work, the passion and the innovation of the original pioneers who paved the way for us.

One more reminder, before we get started. Today's event is being recorded and will be available to all of you in the community after the session. We also plan to take questions through the Q&A button at the bottom of your screen. And we're going to do our best to monitor the chat window so you can also post questions or comments there as well. So, let's get started.

To kick off our session today, I have the distinct pleasure of introducing one of those pioneers, someone that I've gotten to know a bit more, especially through the past year or so, a colleague, a friend and someone that I truly admire. Please join me in welcoming Mr. Dennis Jennings, the former first program director of networking at the National Science Foundation here in the United States.

Dennis Jennings

Evening everybody, thank you very much for the introduction, and thank you especially for all the support from Internet2 for these sessions. Well, good evening, everybody from Dublin, Ireland, as Ana has introduced me, I'm Dennis Jennings, and I'd like to welcome you here to the first Internet Reunion Club session.

The Internet Reunion Club is a fancy name for an excuse for us older, perhaps more charming, more mature, but definitely older Internet guys to get together in person and online from time to time to chat and to reminisce. It's envisaged that each Internet Reunion Club session will start with a public webinar, like this one, at which some aspect of the Internet history will be discussed, followed by a private session for the members of the Reunion Club. Today, we'll be talking about the NSFNET.

The NSFNET, starting in 1985 and winding up in 1995, was an extraordinary achievement, and evolved over that period of time into the global Internet that we know today. We could look at the NSFNET in a variety of ways -- as a dramatic policy, political and financial investment by the US federal government that was successfully transitioned to the private sector, or as an extraordinary technological development as communication speeds went from kilobits to gigabits, or as the

forerunner of the global information infrastructure as search tools and information access tools, including the Worldwide Wide Web, were developed, and definitely as the foundation for the global Internet.

Today, we'll explore this last aspect of the NSFNET, the privatization and commercialization of the Internet infrastructure, as the NSFNET regional networks, part of the original design of the network, evolved into Internet Service Providers, and the US federal government funding for NSFNET was withdrawn. Our speakers today are Glenn Ricart, Elise Gerich, and Susan Estrada.

Glenn Ricart, headed the academic computing and networking at the University of Maryland, and he led the SURAnet regional network project from 1982 to 1993. In that role, he set up the first Internet exchange point in 1988, located in College Park in Maryland, and it connected the original federal TCP/IP networks and the first US commercial and non commercial networks. Later he extended that into MAE East, the Metropolitan Ethernet across Washington. He's also served on boards of public companies, and numerous nonprofit organizations, including the Internet Society and PIR, the Public Interest Registry. Glenn was inducted into the Internet Society Hall of Fame in 2013.

Our second speaker is Elise Gerich, who served as Associate Director for National Networking at the Merit Network in Michigan, and while at Merit she was the principal investigator for NSFNET's T3 backbone project and the Routing Arbiter Project, which we'll hear more about. She has also served as co-chair of the IEPG the Internet Planning Group, chair and co-founder of NANOG, the North American Network Operators Group, National Science Foundation's representative on the federal engineering planning group, and a member of the Internet Architecture Board, the IAB, the top board of the Internet technical structures. Later she served as vice president of the IANA, the Internet Assigned Numbers Authority and its successors from May 2010 to January 2018. Elise was inducted into the ISOC Internet Hall of Fame in 2019.

Our third speaker is Susan Estrada, who founded CERFnet, one of the original NSFNET regional networks in 1988. A commercial network from the get go, as we'll hear, through her leadership and collaboration with PSInet and UUnet, the two big initial commercial networks, Susan helped form the Commercial Internet eXchange, or CIX, enabling the first commercial Internet traffic to be exchanged. Nowadays, she's happily retired, but loves surfing her memory and archives and reliving some of the heady days of Internet growth. Susan was inducted into the Internet Society's Hall of Fame in 2014.

For the panel discussion, and the q&a session, we'll be joined by Guy Almes, the founder and director of the NSFNET regional network SesquiNet, based at Rice University, SesquiNet linked Texas universities and corporates together, and with the backbone of the NSFNET.

The panel itself is chaired by the legendary Dr. Stephen Wolff, one of the many founders and fathers of the Internet. In 1986 Steve became the Division Director for Networking Communications Research and Infrastructure at the National Science Foundation. I had served the previous year in '85 in an earlier role and had made some of the earlier design decisions, but Steve managed and developed the NSFNET project to its extraordinary success, and eventually to its replacement by the global Internet. This included a national backbone network in the US, interconnecting the NSF supercomputers, the NSFNET regional networks, the federal research agency networks, and a whole host of international links to research and education networks. Steve also conceived the Gigabit Testbed, which was a joint project between the NSF and DARPA, designed to prove the feasibility of IP networking at gigabit speeds. And that was a remarkable success. In 1994, Steve left the NSF and joined Cisco and in 2011, he became the Interim Vice President and Chief Technology Officer of Internet2.

So, that's our collection of characters for this evening, and I'll hand you straight over to Glenn Ricart as the first speaker. Glenn?

Glenn Ricart

Thank you so much Dennis. I think that the NSFNET was a truly remarkable achievement that a lot of folks have a lot to do with including you. So you asked me to tell you a little bit about the story of SURAnet. So, what does SURA stand for? SURA is Southeastern University's Research Association Network, and it came about for a couple of reasons. One reason was that SURA wanted to have a supercomputer in the south east, and decided that the best way to have a supercomputer in the south east would be to put in a single south east supercomputer proposal. There was a meeting in October 1983, go and decide which of the institutions in the 18 states and the District of Columbia in the southeast should put together that proposal which they would all back. And the meeting came out with, couldn't decide. However Jesse Poore suggested that what we really needed was a network.

And, by having me not in the room at the time, we managed to get a resolution passed which named me as the person that had to go and make that happen. So, that became the biggest effort at that stage for SURA. And we put together a number of proposals we sent to the National Science Foundation on G Street. There wasn't any Dennis Jennings there yet, our proposals went all the way to the director's office, because they didn't know what to do with them. Never heard of

that kind of a proposal being put together. And in '84 our proposal didn't get any place. In '85 we got a proposal then that finally got awarded in '86. So, that was the beginning of SURAnet.

In terms of commercialization, the first commercial organization on CERN Atlas IBM in Research Triangle, North Carolina. We agreed to put in one SURAnet node in North Carolina, and that node decided that they were going to go and connect IBM. So it's actually up to North Carolina, and what became MCNC network to actually bring in that first of our commercial organizations. And from the very beginning, SURAnet took the opinion that we wanted everyone who could use that kind of a network to be part of the network. So, we had a lot of commercial members.


By the time that we were sold to BBN Planet in 1994, we had several hundred members. I was looking to see if I could find an exact number, but it was certainly over 400 or 500 members that were commercial members by that time. We did that, partly because the NSF kind of forced us to, in the sense, they said we're going to go and withdraw our funding. And we said, hmm, how about we send a bill to everybody of the same size, whether they're academic or commercial, and we'll apply an NSF discount for three years at a reduced rate for the people who are the academic institutions. That way we had preset a price for everybody, and that went forward, and that ended up being pretty successful.

We got started in very late '86, really '87, we started the first connections traffic there. It's about the same time, if I can tell the story of NYSERNet, which was started as a commercial network -- that's Richard Mandelbaum, Marty Schofstell, and Bill Schrader -- and they were going for commercial organizations per se. We were kind of going at SURAnet for commercial organizations, that's kind of a nice way of making all of that happen.

So, I've got a few minutes here. I can show you a couple of slides.

SURAnet Story: “All Scholars in All Disciplines”

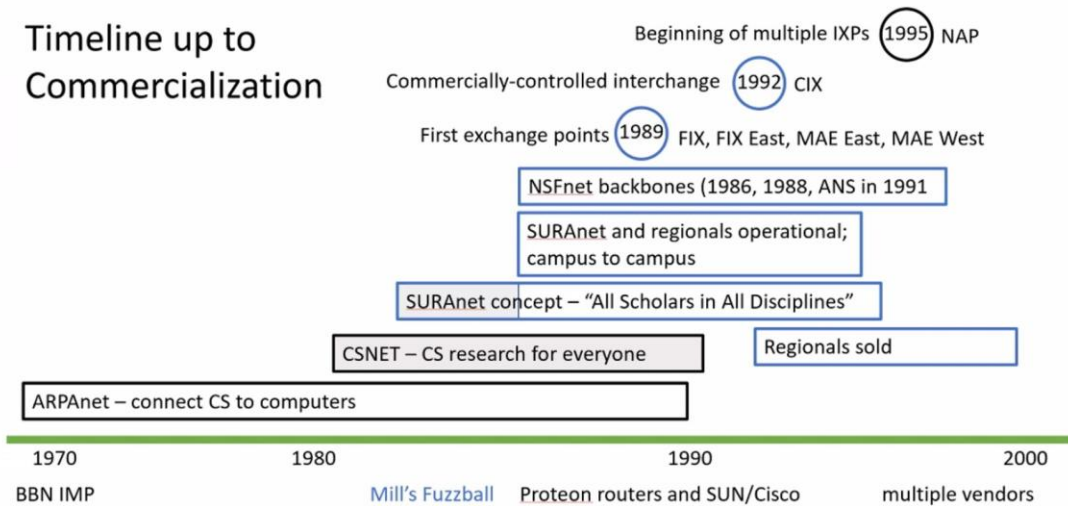
- February 1982: U. Maryland College Park – proposed IRCN based on IP
- October 1983: SURA proposes network for supercomputers
- June 1984: Unsolicited proposal to NSF to create SURAnet – not funded
- 1985: New proposal “U.S. productivity is tightly tied to the ideas that come from university research laboratories. Perhaps the most important way of stimulating new research is to make sure that scientists and engineers in the U.S. can easily and quickly critique and build on each other’s work. A computer communications network is a proven way to create new communities of interest and strengthen those groups which have already formed. In addition, the network can give scientists and engineers high speed, productive access to the limited number of U.S. supercomputers.”
- May 1986: SURAnet Phase I funded (\$932K)
- Dec. 1986-Jan. 1987: SURAnet’s first connection goes live
- 1987: Full phase 1 network operation (13 states + 3 supercomputer centers)
- April 1988: 50 institutions connected to SURAnet
- 1993: Fully self-supporting (end of NSF contributions)
- 1995: 700+ connections; Commercialized: Sold to BBN Planet



Here is a timeline of some of the things I just finished talking about. February '82 we have the College Park proposal to put the campus on an IP network, and the second reason, besides being used for access to supercomputers, is that I had created a campus network which was TCP/IP based, which was a failure because the physics department didn't really want to talk to the business school. The business school wanted to talk to other business schools, the physics department wanted to talk to other physics departments. So, the only way I could correct my mistake was to get other campuses to join the same network, so that the physicists could talk to the physicists, the business schools to the business schools. So SURAnet became all scholars to all disciplines to save me at College Park from having put in a network that no one else on campus wanted to talk to each other for.

You can see here the base I just finished talking about. Phase one network was operational in 13 states plus three supercomputer centers. They had added a number of other connections by the time I said that larger number earlier. We had 50 institutions connected to SURAnet in '88. It was fully self supporting by '93, and we counted almost over 700 connections, here it looks like when I did that counting. It's commercialized and sold to BBN Planet, it instantly became the biggest part of BBN Planet's portfolio. The other folks in that included NYSERNet, which was already part of that, and another network.

Timeline up to Commercialization




I'm showing you the timeline, not the commercialization. ARPANET began in the '70s, CSNET for computer science researchers, the SURAnet concept around the October of '83, until it became operational in early '87. All scholars and all disciplines. The other regionals, NSFNET, really made that NSF wide. NSF wide became government wide. And we started the first exchange point the FIX, which became FIX East, Federal Internet Exchange, and so forth, and that actually did interconnect a lot of traffic which was commercial. So, in '92, the Commercial Internet Exchange broke off that part of the traffic. It was all part of the same layer 2 network, it was running on MAE East, which was Mid Atlantic Ethernet East, and MAE East was really the same layer 2 infrastructure that we were running the FIX on. So, that one interconnection point got logically divided into a commercial exchange and the existing MAE East, in order to go make that happen. And with that, I'd like to turn the story over to Elise Gerich.

Elise Gerich

Thank you, Glenn. Hopefully, Kelly, if you could run my slides for me, I'd appreciate it.

NSFNET: Transition to the Future



Elise Gerich

September 2021

I was invited to speak primarily about the tail end of the NSFNET backbone's cooperative agreement with Merit. So, keep in mind, I'm starting in 1990. And fairly, NSFNET for the T1 backbone had an AUP, an Acceptable Use Policy. A lot of people felt that was incredibly constraining, and that they would like it to be broader. So the AUP was enforced by Merit by asking all the regional networks, that were connected to the NSFNET, to sign a statement that they said that they would be in agreement with the Acceptable Use Policy, and I really felt that that was fairly liberal, because it not only allowed the research and academic community, but it was also allowing for profit firms who were invested in research and education. So, it wasn't as restrictive as just the universities. It did allow outside engagement, and so a lot of regionals agreed to add lots and lots of routes.

Commercial Traffic and the AUP - Early Days

- Regional networks were attracting commercial traffic and felt constrained by the NSFNET's Acceptable Use Policy
- The type of services offered by NSFNET were beginning to be offered by the private sector
- Commercial activity was needed to help maintain the model of a single Internet
- NSF knew when to initiate high technology initiatives and when to get out
- "...if you don't stop old things, then you can't start any new things." Quote from Steve Wolff



So, what I put up here is a picture of the NSF program managers over the 1990 to 1995 timeline. It's Priscilla Huston, George Strawn, Doug Gale, Jane Caviness, and Steve Wolff. Steve, who's our chair for the panel later.

Summer of 1990

- AUP restricted non-research and non-academic use from transiting the NSFNET backbone
- NSFNET backbone was the only high speed game in town
- How to encourage growing use of Internet with fewer restrictions?
- September 1990, NSFNET team announced creation of non-profit Advanced Network & Services



In the summer of 1990, as part of the steps, and as you saw in the last slide, Steve said, you know, once we've been doing something for a while, it's time to try something new. That's paraphrased. So, in 1990, the Merit Cooperative Agreement Group, IBM, MCI, Merit, State of Michigan, and NSF, suggested that we create ANS, which would help to be another step towards the privatizations of

networking, so that NSF could get out of the funding business. In 1990, ANS was announced and the four members there on the screen are Doug van Houweling from Merit, Al Weiss representing ANS, Dick Liebhauer from MCI, and John Armstrong from IBM.

[You are viewing a PDF document](#) [View Options](#)

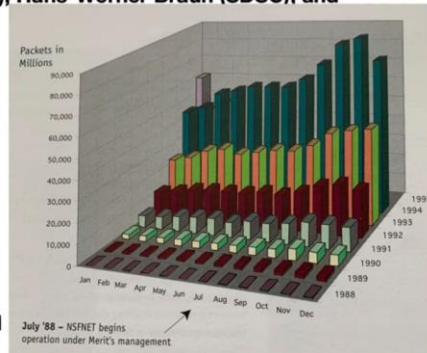
ANS and its subsidiary

- **The NSFNET cooperative agreement was to end in October 1992**
- **NSF extended Merit's cooperative agreement for 18 months (until 1994),so NSF could develop a solicitation**
- **In May 1991, ANS spun off for profit subsidiary ANS CO+RE Systems**
- **ANS CO+RE would attract commercial customers and provide funds to support the backbone**
- **ANS took over the Network Operations Center for the NSFNET Backbone**
- **In 1992, MCI, IBM, Merit and ANS successfully upgraded the backbone from T1 to T3**
- **Merit continues to manage Network Add/Change Requests (NACR) and the Policy Routing Database (PRDB)**
- **The introduction of ANS and Co+RE creates controversy among members of the Internet community**

Basically, this was a step to privatization. It's more than just allowing commercial traffic on the NSFNET backbone, and into the Internet as a whole. Merit and ANS then created ANS CO+RE, which was to be the arm that could actually sell the services that would go over the nationally funded backbone, the NSFNET backbone, and ANS was a subsidiary, or basically it was under contract to Merit. At the same time, because of the increasing requirements of the backbone, trafficwise and everything else, some more technological activities took place, such as upgrading to T3. That included development of new routers, development of new transport systems by MCI, and it also included a new solicitation for a new architecture, so that NSF could start moving out of the business of funding the backbone.

New Solicitation - Summer 1992

- New architecture based on concepts in paper by Bob Aiken (NSF), Hans-Werner Braun (SDSC), and Peter Ford (on loan to NSF from Los Alamos)
- NSF extended Merit's cooperative agreement one last time
- New Architecture with 4 parts:
 - very high-speed backbone (vBNS)
 - Network Access Points (NAPs)
 - Routing Arbiter
 - Regional Network Provider awards
- Award winners announced throughout 1994
 - vBNS to MCI
 - 4 NAP awards - Sprint, MFS Datanet, Ameritech, and Pacific Bell
 - Routing Arbiter - Merit and Information Services Institute (ISI)
- The peering sessions between AS 690 and others were disabled and the NSFNET Backbone was decommissioned on April 30, 1995



In the summer of 1992, the partnership and NSF suggested a new concept, and there was a new solicitation. It was based on a paper written by Bob Aiken, Hans-Werner Braun, and Peter Ford. This included several different steps, it would have a very high speed backbone, it would have network access points, which were basically a lot like the CIX and the FIX, MAE-IX, etc., and it would have a routing arbiter, because one of the things that the NSFNET backbone introduced was the policy routing database, which helped to make sure that routing was not a mess. Because early days, it could be messy, with loops and everything else, and the policy routing database introduced a lot of policies that made sure that you didn't introduce routes to yourself, when you were initiating them-- that's just a very simplification -- and it also provided regional network awards. This was a step towards NSF moving the funding for infrastructure of the the Internet backbone to the private sector. The award winners were: the BBNS was awarded to MCI; there were four network access provider awards, Sprint, MFS Datanet, Ameritech, and Pacbell; and Merit and ISI shared the routing arbiter award.

So, in April of 1995, the NSFNET backbone was dismantled, and this was a big step towards moving towards privatization of the Internet. I know that graph is a bit of a an eye chart, but for those of you who can't tell, it shows that the packet growth, in the years that the NSFNET cooperative project was in place, went from zero packets to over 90,000 packets.

- From: skh
- Date: Mon May 01 14:30:54 1995

With Love To those who nurtured the early Internet

Hi all:


Late last night upon the last hours of NSFNET, I wrote something in response to Sean Doran's poem. By request of a few friends, I am submitting it to NANOG list for thought. I hope you enjoy it.

Pathways of the NSFNET (1988-1995)

All pathways lead but inward, yet ever outward.
All travels begin, and run far only to end at the beginning.
Dreams which lead us to learn to build,
Lead us onward only to find out
How little we truly understand our pathways
which go ever onward in spite of us;
For time burns all but truth away.

Networks begun grow toward success.
If they succeed, they often transform beyond us.
Like children who grow up and beyond us
and lead us further to understand who we are.

(A poem for the end of the NSFNET era from a nsfnert nerd.)




I just wanted to put up this poem that was written by a colleague of mine at Merit as we shut down the network. The picture is of the Merit NSF NET NOC, and this was written by Sue Hares. That's Sue H-a-r-e-s, because there was another Sue Harris, and we don't want to confuse the two, this was by Sue Hares.

And thank you, I'd like to turn this over now to Susan Estrada.

Susan Estrada

Thanks, Elise.



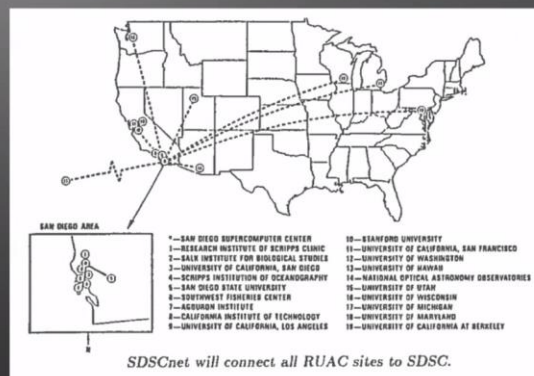
CERFnet and CIX- The Origin Stories

Susan Estrada
September 14, 2021

Thank you. I'm happy to talk about the CERFnet story today, and also the CIX story a little bit, chalking up another victory for truth connectedness in the Internet way.

First, SDSCnet

- RUACs - Remote User Access Centers
- 56 kbps Star satellite network
- 56kbps land connections to locals
- TYMNET @1200 baud (woohoo!)
- ARPANET in Spring 1986
- 56 kbps NSFnet origins to NCAR, NCSA in April 1986



Well, I wanted to talk a little bit about SDSCnet first. The San Diego Supercomputer Center was where I started with this research and education networking. We had actually created a star network in 1996 to all of the universities who were using the Supercomputer Center, or the supercomputer, our Cray. We had a 1200 baud Tymnet connection, we hooked up to the ARPANET

in the spring of '96, and 56 kilobit connections to NCAR and NCSA went in in 1986 as well, those were kind of the precursors to the NSFNET.

So, in June '88, we started writing the proposal for CERFnet. David Walker from UC Irvine and I had attended a workshop at Cornell hosted by the NYSERNet dudes. It was the gathering of the first regionals, and soon to be regionals. On the way home, we decided that Southern California needed to get on board. So, I thought David should organize it, and he thought I should organize it. So, unfortunately for me, he won. With help from David and others, we built the coalition and CERFnet was born. At the same time, you can see on the map, Los Nettos in the LA basin, and led by Jon Postel, was also created, and we decided from the start to cooperate.

In June 1988, a proposal was written...

- 1.544 Mbps screaming fast backbone
 - Full and free transit with Los Nettos
- **33 academic institutions**
- **2 industrial users**
 - SAIC and Northrop
- Always commercially focused:
 - “businesses in the region will be encouraged to affiliate with the network”
- General Atomics/San Diego Supercomputer Center operated
- Funded and live in 1989
- Susan, Susie Arnold and Karen Armstrong were the core implementation team.

CERFnet Topology

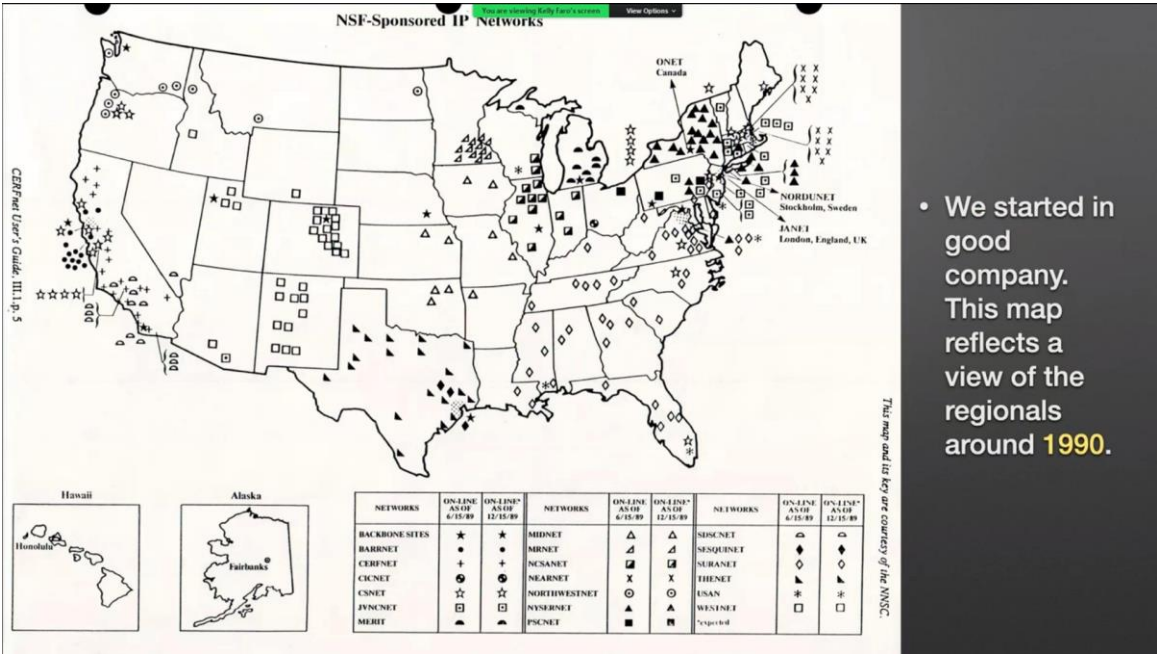
5-28-88

You can see we started off with 33 academic institutions and two industrial users, both SAIC and Northrop. In the proposal, I found the words -- weirdly, I had a copy of the proposal in my pile of crap -- businesses in the region will be encouraged to affiliate with the network. So, we really shot off with the idea that we would use businesses to be part of our network, right from the get go. We actually were up and running before 1989 when we were funding. Me, Susie Arnold, and Karen Armstrong were the core implementation team.




Dedication party
July 10, 1989
(Thanks, Vint!)

Here's a dedication party in July 1989. My newborn daughter was six weeks old, I came to work. I smashed a fake champagne bottle with Vint over a fake Cisco router to christen the new endeavor.



• We started in good company. This map reflects a view of the regionals around 1990.

I found this in the archives too, I think some of you might recognize yourself in this diagram, if that many of you can see yourself, and we had a good and fun company. It was really quite a heady time.



- **1990**
 - CERFnet won the 1990 Interop Achievement Award (thanks, Dan)
 - New service deployed - Dial 'n CERF.
- **1991**
 - 3 international connections went live: Brazil, Mexico and South Korea.
 - 2nd anniversary — the network had grown to over **100 members**.
 - FrEdMail gateway to CERFnet - Internet access to K-12 (Thanks, Yvonne and Al!)
 - CERFnet joined with PSInet and UUnet to form the Commercial Internet Exchange (CIX) interconnecting the networks for commercial Internet service.

In 1990/91 were very busy years of growth and expansion. In 1990, CERFnet won the Interop Achievement Award thanks to Dan -- Hi Dan, I saw you were on -- and we display a new service called Dial 'n CERF, which was the first one of the first dial up connections. In '91 we put in three international connections, to Brazil and Mexico and South Korea. We had our second anniversary, and we had grown to over 100 members by that time from our original 30 plus. A FrEdMail gateway was added to CERFnet as it allowed Internet access to K-12, to Yvonne and Al -- I saw Yvonne was on here too -- and CERFnet joined with PSI and UUnet to form the CIX.

A Word or Two about the Commercial Internet Exchange (CIX)



- The Networks — UUNET, PSINet and CERFnet
- The People — Rick Adams and Andrew Partain, Bill Schrader and Marty Schofstell, Susan Estrada and Susie Arnold
- The Point — Provide neutral connectivity among cooperating carriers with no restrictions on the type of traffic allowed and no settlements (charges based on traffic volume)
- 1991 — CIX was announced in March and CIX-West Santa Clara interconnections up in July.
- Our Customers — Requested this interconnection
- At the time, there was a LOT of NSFnet/NREN drama - ANS, poaching, interconnect fees, settlements, etc.

So, I want to talk a minute about the Commercial Internet Exchange. There were three networks that were involved, UUnet, PSInet and CERFnet. The people were Rick Adams and Andrew Partain from you UUnet, Bill Schrader and Marty Schofstell from PSInet, and me and Susie Arnold from CERFnet, and we all got together in Bill's condo one evening and hammered it out and decided that we were going to go forward with this. The point was to provide a neutral connectivity among cooperating carriers, no restrictions on the type of traffic allowed, and no settlement. So, no charges based on traffic volume. We announced it in 1991 March, and we went live at the CIX West Santa Clara interconnection in July of 1991.


And why did we do it? Because our customers wanted it. At the time, there was a lot of NSFNET/NREN drama, ANS poaching interconnect fee settlements, our customers were getting a little frustrated, as were we, with all of the drama, and we decided this was the quick and dirty way to get past the drama, and move on to actually commercial Internet service.

1992 and 1993

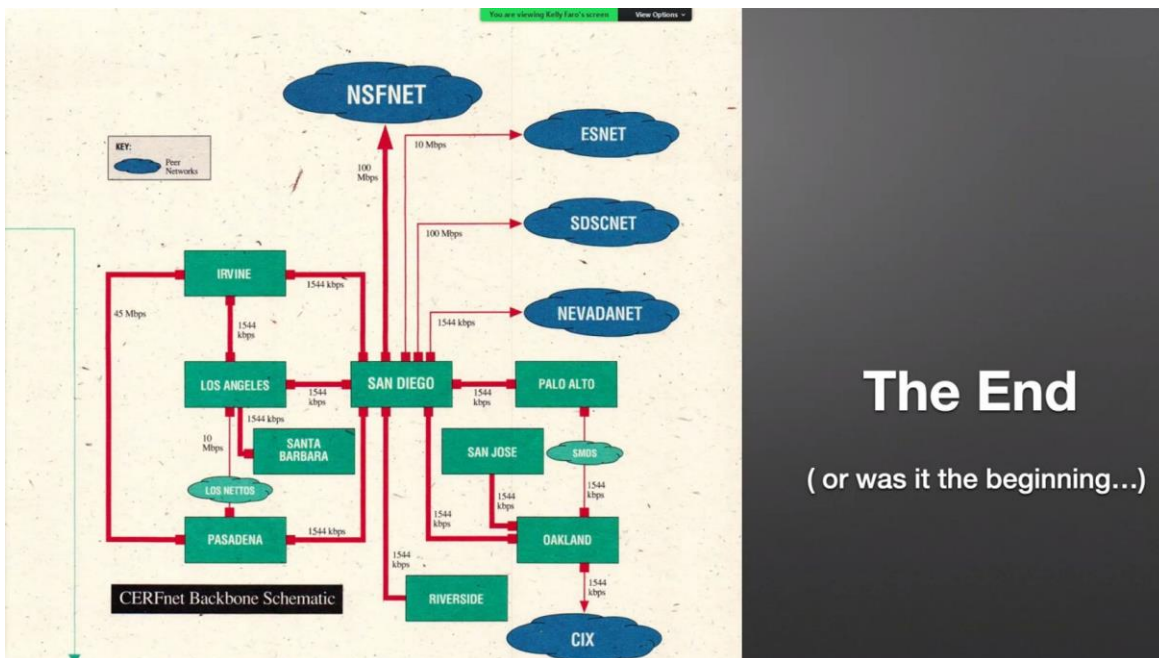
- Commercial traffic allowed
- 24/7 support
- Services had grown to include:
 - CERF 56
 - CERF 1544
 - CERF 45000
 - CERF 14.4
 - Dial n' CERF Plus
 - Dial n' CERF - hourly
 - Dial n' CERF USA - Toll free #
 - Dial n' CERF AYC - unlimited
- CERF n' Safari - CERFnet in the K-12 classroom (Thanks, Yvonne!)



'92 to '93? Commercial traffic was allowed. We had 24/7 support, which we actually had from the beginning because we had the luxury of being in the Supercomputer Center where somebody was there all the time. At this point, though, we now had a T3, we had tons of different dial up services and, luckily for us, we had a great name so we could have clever names for the services. And we did start Cerf n' Safari, which was a CERFnet in the K-12 classroom. Yvonne came to us one summer on an and developed curriculum activities, and we sold that to teachers, so they got the curriculum activities, they got a connection to Dial n' Cerf unlimited for free, and it turned out to be a really great starter for getting K-12 teachers online.



And by 1994, I was no longer with CERFnet, but I did find this map, and you can see that the basic topology concept was still in play, with major universities serving as hubs, and many spokes coming off of them.



It was a very heady time. In this slide you can see peer networks, a solid backbone, and great customers. Thanks for listening, and I'm looking forward to an interesting discussion. Steve and Dennis, I'm turning it back to you guys.

Steve Wolff

I'm here. Thank you, panelists for speaking. I have to say that really what you said brings back a whole host of very dim and, in many cases, I'm sure faulty memories. But one thing that struck me partly from from what you said, and from some of the material that Dennis circulated before this session, was the extent to which the NSFNET started out with basically homemade routers, the fuzzballs for the 56 kilobit network, the giant IBM conglomerations for the T1 NSFNET, but ultimately, networks started using commercial equipment, and I wonder how that transition happened. For example, I know Dennis, when you started JvNCnet, if you were using Vaxen for routers. Sergio -- Hi Sergio -- came in, and I think, what, Cisco? I wonder, for example, how SDSCnet started out, whether they started with with homemade equipment, or immediately were Cisco. Susan, is that something you could answer?

Susan Estrada

Yeah, I know that. We actually -- you know, Sid came from the DOE world -- so, we kind of cannibalized the system that -- I don't know what ES net was called back in the day, but whatever the DOE people used on their network, so it was the same software. It wasn't TCP/IP based, although I know that SDSC was the first site to write TCP/IP code, so that you could talk to the Cray using TCP/IP, but it was called MFE net in the DOE world.

Glenn Ricart

I think they used ANF10.

Elise Gerich

At Merit, they had their own Michigan Control Protocol, and ran that on their own boxes. So, it took a while to transition to TCP/IP.

Steve Wolff

Guy, I haven't heard from you.

Guy Almes

Mm hmm.

Steve Wolff

What were things like in Texas, at Sesquinet and its successors?

Guy Almes

One of the emphases that I would make is that the NSFNET, particularly beginning in 1988, with the beginnings of the T1 based backbone, provided both huge improvements in performance but also great improvements in clarity and simplicity. The so-called Policy Routing Database was actually one of the major contributions that Merit made in that endeavor, both before 1988 and after 1992, we never had such clear routing story to tell, and, of course, the experience of exponential growth in performance is very heady. Going from 56 kilobits to 1.5 megabits, over a period of 12 months or so, was exciting. What it mostly meant for users was that you went from a highly congested 56 kilobits to an uncongested 56 kilobits, and the applications of the day preselected that. But, the thing I'd like to concentrate on is that, for lots of universities that had just not been in the Internet world at all, they were suddenly not only connected, but they were connected in a way that was commensurate with the very high prestige places. So, places like Texas A&M, and even some HBCUs in Texas, were getting performance levels that we're heretofore part of an elite ARPANET group. So, that's one thing I would stress

Glenn Ricart

In SURAnet, we had passed through a phase where we were using the Dave Mills fuzzballs. The reason was that David passed through the University of Maryland on his journeys around, and I'll just show you a quick slide here.


BBN Imps and David Mills' Fuzzballs



Interface Message Processor

\$250,000 one-of-a-kind product of BBN Research

DEC's lowest-end commercial PDP-11; \$3,000 each



He developed -- on the left there is an IMP, Interface Message Processor, the first ones cost a quarter of a million dollars -- the Fuzzball's on the right, it ran on DEC's very lowest end machine, which were about \$3,000 each, which kind of made it a reasonable thing to go do. But, when we started SURAnet, we actually tried with the Proteon routers. Proteon routers were kind of fresh, and it was not always reliable, but we found out that if we would unplug them and replug them, they could come back up and work more efficiently again if they had failed. So, that whole reboot your machine is something that we were doing a lot with the original Proteons. Of course, the software got better and better, and that became much more reliable.

In terms of commercialization, I might also mention that one of the impacts of the NSFNET regional networks was that the early router vendors, including Proteon, and especially Cisco, suddenly had customers, and it was interesting when you could just pick up the phone, call the Cisco phone number and get the CEO, and place an order for a router. That was very, very interesting. It didn't last very long, but it was exciting. So, in a sense, the commercialization of the Internet was happening, in a sense, even before you started to have commercial sites connect.

Some commercial sites did come on fairly early, but I want to tell you a little Cisco story, pre-Cisco, there was Sun Microsystems, and we heard that they had a good router. And of course, the Proteons were having the problems I just talked to you about, so we said, Hey, how about getting some SUNs in here to be routers for us? They said, You're too far away, we can't deal with that. You know, we might sell you some machines later. So, we ended up sticking with Proteon.

Elise Gerich

So, a lot of folks ran gated, didn't they? On just Unix boxes, or am I off the map here?

Guy Almes

If you had a Unix box, then gated was very important.

Steve Wolff

But routing's always been a problem. Several of you have mentioned the Policy Routing Database, but one of the lasting contributions of the NSFNET program, viewed in different lights by different people, is BGP..

Guy Almes

Uh, huh.

Steve Wolff

...and I wonder how BGP made its way into the regional networks?

Elise Gerich

So, can I jump in with a story? I know we were running BGP on the NSFNET backbone, and I think Cisco was the only vendor that had implemented it to begin with, but many of the regionals were running routers that had both BGP and EGP as their protocol, and we were supporting both from the NSFNET perspective, and we were having difficulty getting everyone to take the leap, to jump with faith, that if they turned on BGP, and shut off their EGP, that they wouldn't crash and burn. So, this is the great thing about a research and academic network, is sometimes you can just do stuff, and so one night, we turned off PSC's, Pittsburgh Supercomputer Center's, non-BGP routing announcements, and so they ran for all night and most of the next day, just using BGP routes and, oh my gosh, everybody sort of fell over and started using BGP. But, it wasn't like we had an agreement, I think we just sort of decided to give it a try, and if it broke, we could always bring up the other connection. That's my recollection.

Guy Almes

If I can speak of the pre-NSFNET backbone era, Elise, before the T1 backbone, there's some fascinating stories dealing with the interconnection of NYSERNet and SURAnet who, as it happened, we're both running the RIP interior routing protocol, and for reasons that, at the time, no one understood, it worked in ways, as I say, in ways that were beyond what was intended by the design, strange interactions between the RIP protocol and its notion that 16 was infinity, and notions of default routes, made things work in ways that only, months later, they understood why they were working.

Guy Almes

Then, as I say, during the backbone, not only was there great performance, but the regional backbone architecture more or less allowed most people to use default routes most of the time. In a sense that was a world in which EGP was adequate, whereas before the backbone EGP was really failing badly, and it's important that BGP did arise by the early 90s, when unanticipated patterns of connection were being built.

Dennis Jennings

Dennis here, can I jump in? I'm curious as to how important Interop was in this whole saga, in the development of the product side, the commercial product side? Do you mind if I ask that question of the panel?

Steve Wolff

You've done it.

Dennis Jennings

Indeed.

Steve Wolff

Panelists?

Glenn Ricart

I'll jump in, and so, let's see, the first Interop was -- Dan, are you around? -- it was '86, was the first Interop? Let me pretend that it was. The year before, I had called a meeting of people who were working on stuff to just get together, and try to make sure we were all doing some of the same things, and it was the first time I'd ever called a meeting using email and everyone came, they didn't even have to have a letter that had been signed, without the physical mail in order to get people to come to a meeting, and Dan was one of the folks that was there. They said, well, are you planning on holding this meeting again next year? And I said, no, and he says, Well, I'd like to hold this meeting, and I said, Great, we'll send everyone to your meeting, and so we got also really good attendance at the first Interop.

Elise Gerich

Susan, did you want to jump in?

Susan Estrada

Yeah. I'm just reading the chat, and people on the chat know more about this than any of us do, for sure. So '86, in fact, was the correct year, and I kind of remember that meeting in Monterey, and I have to say that Interop, for me and for our network, was a tremendous opportunity, because, number one, a lot of people didn't know we were out there, because we were insulated in our R&E world, and number two, it was an opportunity for us to speak to a lot of the vendors who were there, and vice versa. So, I don't know how we would have gotten past a lot of the issues that we were running into, you know, over the time, if Interop wasn't there to sort of stir the pot -- the witch's brew, Dan. I was reading through some of the archive information I have, and in 1992, one of the big things for us was to try to convince people that the Internet was a cool thing, and they should use it, and now we worry about where we should put our selfie mirror, so we've really come quite a distance, and I think Interop was kind of a pivotal place for us all to get together and yak about this stuff in a different way than we would have had any other way.

Glenn Ricart

And it was the only place that you could go and try out different pieces of equipment from different folks, from different software, and see if it actually ran. So, it was really important at the beginning, because of the true interoperability, the plug tests to go and make sure that things would work together, and it formed a very important function from that point of view.

Elise Gerich

I agree, I think Interop was a key, often overlooked, activity that happened and supported the growth of the Internet, by allowing there to be non-proprietary solutions, because all the equipment had to talk to each other. If you looked at other backbones before, then you had to buy vendor A because you wanted to talk to your other equipment that only spoke vendor A's proprietary software, and so Interop allowed that interoperability -- that's why it has that name -- and it moved us from proprietary solutions to ones that were able to be explored across the world, in different kinds of platforms, it was really a great thing to help us grow.

Steve Wolff

I certainly would agree that Interop was very important in the early days.

Steve Wolff

I have another question. Having to do with the commercialization and privatization of the NSFNET. Of course, NSF had a fairly strict acceptable use policy, which was dictated by its advisory structure, but, as you all know, it was enforced more in the breach than in the observance, although, in Washington, what do we know? The contact between NSF, NSF's superstructure, and the regionals, was basically through the program directors for networking, who you've already heard -- Priscilla Huston, Jane Caviness, George Strawn, and Doug Gale -- and I'm wondering, they were sort of given the instructions not to be too strict about the AUP when they talk to the regionals and universities, and I was wondering how -- if any of you had any experiences with discussing that particular topic with those particular people?

Guy Almes

It's actually an anecdote that relates to you, Steve. Who was the wonderful executive director of FARNT for a period of time? I forgot, I've blanked on her name, but..

Guy Almes

Laura Breeden.

Guy Almes

Laura Breeden had a conversation, I believe with Steve. The question was something like, suppose that a company is connected to the Internet in one place, and suppose a university is connected to the Internet in another, and the university has a professor who wants to buy test tubes from the company. Can they use the NSFNET to order the test tubes? And the answer was yes. And Laura was pleasantly shocked. Shocked, I say. So, I think this relates to what you were saying about, your primary energy was not going into enforcement of the AUP, and yet, as it became evident that the Internet was really useful for real work, and not just for tiny things, the companies did get to have bona fide commercial connections to the Internet, and that led to the growth of the commercial Internet, which continues to this day.

Glenn Ricart

We told all our folks at SURAnet about the one end rule, which required one end to be academic or research connection, and that solved most of the problems, and when commercial organizations talked to other commercial organizations, we didn't necessarily know that was happening. If anybody asked us we pointed out that mightn't be compliant with rules, but we didn't really have a way of telling if anyone on SURAnet was doing something that may not have followed all the rules.

Guy Almes

There really is another important driver of the commercial Internet that is often not pointed out, and in addition to the routers and the equipment, which is one story, through Interop and so forth, and also the eventually commercial connection as to companies, there is a third thing that was going on, and that is that, in classic NSFNET, the sites connected were universities and colleges, and initially, that meant computer scientists or supercomputer users. Eventually, it meant students and faculty of all departments, which eventually became enormously useful. But, it also percolated to the dormitories, and it's important to note that the American university is a strange institution, in that every spring, a university gets rid of its 25% smartest people, and they go out into the world, and so think for a moment about the undergraduate students who had been exposed to the NSFNET in the early 1990s, who graduated, then they went to join companies, and explained to the companies how great the Internet was, and they also went to apartments and homes and started finding out how do I get on residential Internet connections, and so both of those rules, those students that we graduated during the early 90s, were really the unsung, how to use Vint Cerf 's job title, Internet Evangelists, and those people were hugely important in driving the exponential growth of the commercial Internet during the late 1990s.

Susan Estrada

Just a quick thing, we never really enforced the AUP. Everybody that came on our network, was told that they were supposed to use it for research and education purposes but, obviously, some of the companies could speak to each other and still do, the commercial companies could still follow under that criteria. I used to spend some time looking at the routes and stuff to see who was talking to whom, just because I'm nosy, and it was fascinating to see how much was going on between the commercial companies, and not just between the researchers, or within the university systems. To Guy's point too, I had one customer who -- I think he started off at Hughes-Peters -- and he changed jobs every now and then, and every time he changed jobs, he came back on and got a new connection for his new company. So, you know, once you're spoiled, you got to have your Internet.

Elise Gerich

Yes. So, from my perspective, I know, at Merit running the NSFNET, we had to say to each regional that if you're going to submit routes, and I think they signed a blanket agreement with us, that then I shipped off to NSF, to Steve, that said that you would only submit routes that were in support of the general principles of NSF AUP, and we didn't have a lot of AI stuff then, so we weren't delving into packets to see if anybody was being bad. We weren't analyzing the data that went through the connections, we were only looking at, okay, everybody's saying that they're being a good citizen, and that they're using this for the general purposes of the NSF AUP, and I think that was a great way to do it, because this was seeding research and development in network and Internet infrastructure. A lot of commercial things came out of the NSFNET activity and their experiment, and NSF, like my slide said the quote from Steve, I can't remember it exactly but, Once you've done something, it's time to try something new, and that's what NSF did. They let us try this, it grew, and then they decided, Okay, now to try something new, and they moved on, and you guys had already figured out from the regionals, how to get funding to continue to do your activities that you needed to do. So, I thought NSF was very liberal, though there were disagreements from some members of the community who wanted to start their own enterprises, and we certainly at the NSFNET competed with that.

Dennis Jennings

I was going to ask that question. There are obvious tensions that must arise with this Nelsonic telescope to the blind eye. When the commercials began to get some muscle, they must have started complaining, and say, Hey guys, you're using federal money to undermine me? When did the first lawsuit fly? And what was it all about? I mean, this is a natural progression. Does anybody here know?

Susan Estrada

You know, I was talking to Bill Schrader yesterday, as a matter of fact, and he said that he had written a letter to the NSF, which Steve may remember, or might have chosen to forget about, in around 1995, saying that it was not fair for the university folks, or the the National Science Foundation, the government, to be competing with commercial networks, and that was about the time when the NSFNET was looked at, and eventually shut down for those reasons, but that's just hearsay. But Steve would have a better recollection, or perhaps deny.

Steve Wolff

In a sense, it all became moot in 1992 with the passage of the Boucher Amendment, PL 102, mumble mumble, which said it was perfectly okay for commercial traffic to use the NSFNET provided it didn't materially hinder the R&D traffic. But yes, a lot of that came about because of the quite natural competitive concerns of people like Bill Schrader and Rick Adams, and others. I think probably, in my recollection, a good bit of my last year at NSF was spent answering various Freedom of Information Act requests from Gordon Cook, and others, and coping with the IG investigation, which went on. One of the organizations that played a major role in the development of the Internet, that hasn't been mentioned here, is the IETF. I was wondering if -- to what extent the regionals had interactions with the IETF, attended IETF meetings, kept abreast of the RFCs, and so on.

Guy Almes

Yes, very much.

Steve Wolff

Any comments on that?

Glenn Ricart

Very much.

Guy Almes

In fact there was a period of time when I chaired a so-called Interconnectivity Working Group, which had to do with sorting out the problems that were latent in the EGP protocol as the topology of the Internet became more interesting, and the primary concrete deliverable was an unintended one. We had strict marching orders not to create a new protocol, but the primary thing that came out of our working group was the creation of BGP.

Glenn Ricart

Yeah, the regionals had to stay very much up with the IETF, and we sent lots of folks to the IETF, until commercial equipment became more stable, because in the beginning, the commercial equipment was not quite as stable, you had to understand what was going on, why was doing this, what standard was it operating with, why wasn't it doing this other thing. By the time that Cisco and other folks came up with much more reliable routers, that decreased the need for the regionals to be attending the IETF. But I have to say that a lot of folks who started, continued, they just weren't doing it because it's under regional cover, but they stayed involved in the Internet community.

Elise Gerich

I know Merit was very engaged at the IETF, not only through the NSFNET project, but the Merit engineers, Mark Knopper and others, were quite engaged in the IGP protocol wars in IETF, and that was something that they continued to be active in.

Susan Estrada

Just one quick scape. I was on the IESG for a couple of years in the early 90s, I guess, and I was one of the co-chairs of the Operations Group, and it was an interesting thing, because the IETF was really looking at the future, and operations looking at, Holy shit, what just happened? And so, we had a really hard time kind of trying to figure out the point of operations within the IETF.

Steve Wolff

Thank you all. I mean, I'm being carded by Ana that our time is up. So, Ana, take it away, and thank you.

Susan Estrada

Thank you.

Ana Hunsinger

And apologies, I have to jump in, but we are at time. So, please join me in thanking Glen, Elise, Susan, Guy, and Stephen Dennis, for joining us today, and for lending a little bit of your time and sharing the stories. Thank you so much. Thank you also to all of you for joining, and for the great participation in the chat window, and your comments, and please keep them going. We hope to host another one of these in a few months, so stay posted for more information on that, and again, a recording of this session is going to be available through the Internet2 website, in the Internet2 Community Voices web page. Please do us a favor and take just two minutes of your

time to fill out the survey that's going to appear after you leave the session. Thank you, everybody, and have a great day.

Dennis Jennings

Thank you, panel. Thank you all, and good evening.

Ana Hunsinger

Thank you all